

SCHEDULER FOR A PACKET ROUTING AND SWITCHING SYSTEM

ABSTRACT OF THE DISCLOSURE

A scheduler allocates service to enqueued cells of connections provisioned for guaranteed service levels employing a structure with one or more of the following features to achieve efficient high-speed packet switching (cell relay). For very high-speed switching of connections, such switches operating up to 10, or even 100, Tbps, burst scheduling of cells is employed in which a number of cells, termed a burst, are serviced when a queue is eligible for service. When a queue has less than the number of cells in the burst (termed a short burst), the scheduler still schedules service, but accounts for saved service time (or bandwidth) of the short burst via queue length when the queue's eligibility is next considered for service. In addition, high and low bandwidth connections of a queue may be allocated into two sub-queues, with priority assigned to the two queues and delay-sensitive traffic (high bandwidth connections) assigned to the higher priority sub-queue. High-bandwidth queues may be assigned high and super-high priority to account for ties in scheduling decisions for eligible queues. The scheduler may tentatively schedule a queue as eligible for service, but employ information generated from output ports indicating an output port's congestion to decide whether a bid for service of the eligible queue should be generated. The scheduler may allocate service to at least two different traffic types separately, e.g., standard traffic queues and guaranteed bandwidth traffic queues, based on traffic type specific algorithms that determine an eligibility for a queue associated with each type of traffic. When eligibility of queues for each type of traffic is determined, then another scheduling algorithm allocates service to each eligible queue of the traffic types based on, e.g., a rate proportional server algorithm.